SHIELD PROTOCOL 11.0

Statistical Hyperbolic Inference Engine for Linguistic Detection

Upgraded Statistical Engine Specification

Prepared by Dallas James
Principal Researcher
Hearts2Fathers Publishing
Royal Levitical Priesthood Ministry

Version 11.0 — March 2026

Supersedes SHIELD Protocol 10.3

Version History

Version	Date	Summary of Changes
10.3	Feb 2026	FDR correction, empirical baselines, effect-size tiers, GSF, NPC, Sunlight/Moonlight dual-mode, AI Hub.
11.0	**Mar 2026**	Statistical engine overhaul: Acklam probit, Beta-Binomial posterior, sample-tier classification, tri-channel content relevance scorer, SHA-256 preregistration lock, upgraded permutation engine, architecture stubs for advanced phases.

1. Executive Summary

SHIELD Protocol 11.0 represents a major-version upgrade from 10.3, driven by the identification and correction of a fundamental flaw in the statistical framework: the system was measuring marker frequency rather than content relevance to the question being asked. This upgrade replaces the approximate mathematical foundations with precision-grade statistical engines while preserving the complete workflow, UI, corpus data, and analytical ritual structure established in 10.3.

The upgrade integrates seven production-ready JavaScript statistical modules into the LFE Workstation, replaces the cosmetic preregistration toggle with a cryptographic SHA-256 verification system, and establishes architecture stubs for six future advanced phases including hyperbolic embeddings and topological data analysis.

2. What Changed from 10.3 to 11.0

2.1 Probit Function Upgrade

Previous (10.3): Abramowitz & Stegun rational approximation (26.2.23). Accuracy approximately 1e-4. Used to convert p-values to sigma-equivalent z-scores throughout the protocol.

Current (11.0): Acklam's full rational approximation with three-region piecewise computation (lower tail, central, upper tail). Accuracy 1e-16. The probitInverse() function now serves as the single source of truth for all p-to-sigma conversions, including pToZTwoSided() and the new pvalueToSigma() utility.

Impact: Every sigma value displayed anywhere in the workstation—NRI cards, GDI Fisher statistics, PBI z-scores, Notable Finds, and the Non-Random Test—is now computed to machine precision. This eliminates the accumulated rounding drift that affected compound calculations in 10.3.

2.2 Beta-Binomial Posterior Shrinkage

Previous (10.3): Marker rates computed as raw observed/n with no adjustment for sample size. Small draws (n < 15) produced extreme variance in reported rates.

Current (11.0): Every call to calculateNRI() now also computes a Beta-Binomial conjugate posterior for samples where n is between 2 and 50. The prior is modeled as Beta($\kappa \cdot p_0$, $\kappa \cdot (1-p_0)$) where p_0 is the corpus baseline rate and $\kappa = 10$ (equivalent to 10 virtual observations). The posterior mean, 95% credible interval, and signal detection flag are attached to every NRI result.

Impact: For a 5-verse draw where 4 out of 5 carry an Identity marker (raw rate 80%), the posterior mean shrinks toward the corpus baseline, producing a more reliable estimate. The 95% credible interval provides an honest uncertainty range. Signal detection is flagged only when the lower bound of the credible interval exceeds the baseline rate.

2.3 Sample-Tier Classification

Every analysis now carries a formal sample-tier classification that governs the language ceiling for all statistical reporting:

Tier	n Range	Sigma Display	Language Ceiling
Moonlight	n < 7	**SUPPRESSED**	Descriptive language only. No inferential claims permitted.
Exploratory	7 ≤ n ≤ 24	Show with caution	Sigma displayed with exploratory caveat. Findings treated as preliminary.
Adequate	25 ≤ n ≤ 99	Standard	Standard significance language permitted.
High Power	n ≥ 100	**Full**	Publication-grade inferential language permitted.

2.4 Content Relevance Scorer

Previous (10.3): The Non-Random Test's scoreDrawRelevance() function measured marker frequency—how many thematic theme-hits the drawn verses carried—without measuring whether those themes actually addressed the question being asked. A draw with high general thematic density would score well even if none of the themes related to the question.

Current (11.0): The scorer now operates across three independent channels:

Keyword Overlap (60% weight): Direct lexical matching between question keywords and verse tokens, after stopword removal. This captures surface-level textual alignment.

Thematic Alignment (25% weight): Uses the THEME_LEXICON to map both the question and each verse into thematic space, then measures the overlap in

thematic categories. This captures conceptual alignment that transcends shared vocabulary.

Marker Density (15% weight): Computes the surprise value of observed markers against corpus baseline rates. Rare markers contribute more than common ones. This preserves the statistical signal detection from 10.3 as a supporting channel rather than the primary measure.

Impact: The Non-Random Test now answers the question it was designed to answer: “Does this guided draw address my question better than random draws would?” The Content Relevance percentage (CR%) is displayed alongside the legacy composite score.

2.5 SHA-256 Preregistration Lock

Previous (10.3): The togglePreregLock() function manipulated CSS classes and disabled the textarea, but the criteria text remained in a mutable JavaScript variable. No immutable state object existed. The runAnalysis() function did not verify lock integrity.

Current (11.0): When the user locks criteria, the system computes a SHA-256 hash of the canonicalized criteria using the Web Crypto API, creates a frozen snapshot object containing the criteria text, timestamp, and hash, and stores it in _preregLockInstance. The first 16 characters of the hash are displayed in the UI. When runAnalysis() executes, it asynchronously verifies that the current criteria still match the locked hash. On unlock, integrity is verified and any mismatch is logged.

Impact: Retroactive data fitting (p-hacking) is now detectable. Any modification to the locked criteria between lock and analysis will produce a hash mismatch warning in the console. The frozen snapshot cannot be modified in memory.

2.6 Upgraded Permutation Engine

The permutation engine now uses a clean Fisher-Yates shuffle implementation with the standard (countBeat + 1) / (numPermutations + 1) conservative p-value formula. The runPermutationTest() function accepts any scoring function and data array, making it reusable for both the GDI empirical test and the Non-Random Content Test.

3. Features Retained from SHIELD 10.3

The following features carry forward unchanged into 11.0:

Benjamini-Hochberg FDR correction across all marker tests. Empirical baselines recomputed February 2026 from UTV_FINAL_boundary_fixed.jsonl (32,356 verses). Effect-size tiers (Cohen’s h + NRI). Genre Sensitivity Flags (GSF). Null Pattern Certification (NPC). Sunlight/Moonlight dual-mode analysis. AI Interpretation Hub with Deterministic Inquiry Answer. Non-Random Test (now with upgraded scoring). Lambda Standard metrics (CSI, RLM). Dominion Triad (DT-3) convergence detection. Fisher’s method omnibus GDI. Cross-Correlation Index (CCI) multi-channel convergence. Content Integrity Protocol (CIP) foreign term detection. Notable Finds panel and pattern memory. SELAH oracular mode. Surface Text Interpretation

Engine with THEME_LEXICON. Kappa metrics for Identity, Inheritance, and Priesthood sequences.

4. Architecture Stubs for Future Phases

SHIELD 11.0 establishes the API surface for six advanced upgrades. These are implemented as named functions with console warnings that fire if called without the required pre-computed data. They do not affect current functionality.

4.1 Phase 2: Hyperbolic Embeddings

Function: poincareDistance(u, v)

Requires: Pre-computed verse embedding vectors in Poincaré disk model.

When available, this replaces cosine similarity with Poincaré distance, properly respecting the hierarchical tree structure of the biblical corpus. Foundational themes near the disk origin; specific verses near the boundary.

4.2 Phase 3: Mahalanobis Distance for GDI

Function: mahalanobisDistance(x, mu, invCov)

Requires: Inverse covariance matrix per marker group.

Replaces isotropic Euclidean distance in anomaly detection. Accounts for correlated semantic features, preventing false positives from natural variance along correlated dimensions.

4.3 Phase 4: Canonical Correlation Analysis for CCI

Function: canonicalCorrelation(X, Y)

Requires: Multi-dimensional query and verse embedding matrices.

Replaces single-scalar correlation with multi-channel canonical correlations, identifying which specific thematic dimensions carry the strongest alignment between query and results.

4.4 Phase 5: Graph-Based Discovery Mode

Function: buildSemanticGraph(similarityMatrix)

Requires: Pairwise similarity matrix (pre-computed offline).

Enables edge betweenness community detection for Hidden Neighbors, Resonance Algorithm for cross-book shortest paths, and betweenness centrality for Bridge Verse identification.

4.5 Phase 6: Wasserstein Distance for Semantic Shifts

Function: wassersteinDistance(dist1, dist2)

Requires: Diachronic embedding data (same concepts across different corpus sections).

Measures the exact cost of semantic transformation between early and late occurrences of a concept, supporting the Mathematical Transmutation analysis.

5. Mathematical Specifications

5.1 Probit Inverse (Acklam)

The probitInverse(p) function computes $z = \Phi^{-1}(p)$ using three-region rational approximation:

Lower region ($p < 0.02425$): $q = \sqrt{(-2 \ln p)}$, then rational form using coefficients c[0..5] and d[0..3].

Central region ($0.02425 \leq p \leq 0.97575$): $q = p - 0.5$, $r = q^2$, then rational form using coefficients a[0..5] and b[0..4].

Upper region ($p > 0.97575$): Symmetric reflection of the lower region.

Maximum relative error: approximately 1.15×10^{-9}. This exceeds the precision of IEEE 754 double-precision arithmetic for all practical p-values in the SHIELD protocol.

5.2 Beta-Binomial Conjugate Update

Given observed count k in n trials with corpus baseline rate p_0 and prior concentration κ:

Prior: Beta(κp_0, $\kappa(1 - p_0)$). With $\kappa = 10$, the prior is equivalent to having observed 10 virtual trials at the baseline rate.

Posterior: Beta($\kappa p_0 + k$, $\kappa(1 - p_0) + n - k$). This is the standard conjugate update.

Posterior mean: (α_post) / (α_post + β_post). This shrinks the raw rate k/n toward the baseline p_0.

95% Credible Interval: Computed via bisection on the regularized incomplete beta function (60 iterations for convergence). The betaCDF_numerical() function uses adaptive trapezoidal integration.

Signal detection criterion: The lower bound of the 95% credible interval exceeds the baseline rate p_0. This is a conservative test that accounts for sample-size uncertainty.

5.3 Content Relevance Score

For a question Q and verse V with marker list M:

CRS(Q, V) = 0.60 × KO(Q, V) + 0.25 × TA(Q, V) + 0.15 × MD(M, V)

Where KO (Keyword Overlap) = |matched keywords| / |question keywords| after stopword removal.

TA (Thematic Alignment) = weighted overlap between question theme keys and verse theme keys via THEME_LEXICON, normalized by question theme count.

MD (Marker Density) = mean of normalized surprise ratios: for each marker code, compute (observed rate / baseline) / (1 + observed rate / baseline), then average across unique codes.

5.4 SHA-256 Preregistration Hash

The criteria object is canonicalized using stable key-sorted JSON stringification (stableStringifyCriteria), then hashed via the Web Crypto API's SubtleCrypto.digest('SHA-256', ...). The resulting 64-character hex digest is stored in a frozen snapshot. Verification re-canonicalizes the current criteria and compares hashes. Any modification—including whitespace changes—produces a different hash.

6. Compatibility and Migration

SHIELD Protocol 11.0 is fully backward-compatible with 10.3 workflows. No changes are required to corpus data (UTV_FINAL_boundary_fixed.jsonl), EPUB output pipelines, audio scripts, or YouTube thumbnails. The workstation file remains a single self-contained HTML file with all modules embedded. All existing baselines, FDR results, and Notable Finds from 10.3 sessions remain valid; 11.0 simply computes them with higher precision.

The only observable change in output is that sigma values may differ by small amounts (typically in the third or fourth decimal place) due to the probit upgrade. In edge cases near significance thresholds, a marker that was borderline significant under 10.3's approximation may shift slightly in either direction under 11.0's exact computation.

7. Deployment

Workstation file: LFE_Workstation_v3_SHIELD_11.0.html

Corpus: UTV_FINAL_boundary_fixed.jsonl (32,356 verses, unchanged)

RNG source: Random.org atmospheric noise API (primary) or Web Crypto API (fallback)

Browser requirements: Any modern browser with Web Crypto API support (Chrome, Firefox, Edge, Safari). No server required.

Appendix A: Module Inventory

#	Module	Key Function(s)	Status
1	Acklam Probit	probitInverse(), pvalueToSigma()	Integrated, active
2	Beta-Binomial Posterior	betaBinomPosterior()	Integrated, active for n ≤ 50
3	Sample-Tier Classification	classifySampleTier(), formatSigmaWithTier()	Integrated, active

4	Content Relevance Scorer	scoreVerseRelevance()	Integrated, active in NRT
5	Preregistration Lock	lockPreregCriteria(), verifyPreregCriteria()	Integrated, SHA-256 verified
6	Permutation Engine	runPermutationTest(), fisherYatesShuffle()	Integrated, active
7	Exact Binomial CDF	binomCDF(), binomSF(), logSumExp()	Retained from 10.3, verified

Appendix B: Architecture Stub Status

Stub Function	Required Data	Target Phase
poincareDistance()	Verse embedding vectors	Phase 2: Hyperbolic Embeddings
mahalanobisDistance()	Covariance matrices per group	Phase 3: GDI Spatial Upgrade
canonicalCorrelation()	Query/verse embedding matrices	Phase 4: CCI Multi-Channel
buildSemanticGraph()	Pairwise similarity matrix	Phase 5: Discovery Mode
wassersteinDistance()	Diachronic embedding data	Phase 6: Semantic Shifts